arXiv:1212.1320v1 [math.DS] 6 Dec 2012

Complexity as a homeomorphism invariant for tiling spaces

Antoine Julien
Department of Mathematical Sciences
Norwegian University of Science and Technology (NTNU)
N-7034 Trondheim, Norway

Abstract

It is proved that whenever two aperiodic repetitive tilings with finite local complexity have homeomorphic tiling spaces, their associated complexity functions are asymptotically equivalent (it implies, if the complexity is polynomial, that the exponent of the leading term is preserved by homeomorphism). This applies to d-dimensional infinite words: if two $\mathbb{Z}^d$-subshifts (with the same conditions as above) are flow equivalent, their complexity functions are equivalent. How this result relates to the theory of tilings deformations is addressed in the last part.

1 Introduction

A $\mathbb{Z}^d$-subshift is a closed, shift-invariant subset $\Xi \subset \mathcal{A}^{\mathbb{Z}^d}$, where $\mathcal{A}$ is a finite set (alphabet). An important tool in the study of subshifts is the complexity function. Given an element $w \in \Xi$, its complexity function $p_w(n)$ counts the number of subwords of w of size n. In dimension one, the definition is clear; in higher dimension, there are several definitions for the complexity function, for example $p_w(n)$ can count the number of cubic subwords, *i.e.* subwords which are included in a cube of edge length n. The complexity functions considered in this article will always be equivalent to this definition.

The complexity is related with some important combinatorial properties of the word. For example, if $p_w(n)$ does not grow at least linearly, then w has d independent periods, and the subshift is finite (this is due to Morse and Hedlund [14] for words of dimension one; see also Lagarias and Pleasants [13] for an analogue of this theorem in the more general case of Delaunay sets of $\mathbb{R}^d$).

The point of view which is adopted in this article is that of tilings: a tiling is a covering of the plane by geometric shapes (tiles), with no holes and no overlap. A d-dimensional infinite word defines naturally a tiling of $\mathbb{R}^d$ by cubes of different colors. A tiling can be associated with a $\mathbb{R}^d$-dynamical system in the same way a word is associated with a $\mathbb{Z}^d$-subshift. In the case the tiling is a tiling by cubes, its tiling space is called the *suspension* of the subshift. Using the tiling perspective allows for greater generality.

It is assumed that the tiling is made using finitely many tiles up to translation, and that there are finitely many ways to fit tiles together. This *finite local complexity* condition is necessary to define a complexity function for tilings. In this case, $p_T(r)$ counts the number of patches in T, of radius at most r. Note that there are also ways to define a complexity function in more sophisticated ways for tilings of infinite local complexity [16].

The main result of this paper is that two homeomorphic tiling spaces (with some additional assumptions on the tilings) have asymptotically equivalent complexity functions. This statement deserves a bit more precision.

Two complexity functions p_1 and p_2 are said to be equivalent if there are some positive constants m, M, C_1, C_2, such that for all r big enough,

$$C_1 p_2(mr) \leq p_1(r) \leq C_2 p_2(Mr).$$

A few remarks:

- If the functions p_i are at least linear (which is the case for non-periodic tilings, by the Morse–Hedlund theorem), then C_1 and C_2 can be taken equal to 1 (up to adjusting the values of m and M);
- On the opposite, if the functions p_i grow at most polynomially, then m and M can be chosen equal to 1 (up to adjusting C_1 and C_2). In particular, if p_2 grows like a polynomial, then so does p_1, and the exponent of the leading term is the same.

The first observation (Proposition 2.8) is that any two reasonable definitions of complexity function for a tiling are equivalent in the sense above (namely, remark that taking any two norm for defining "patch of size n" would give equivalent complexity functions—all things being equal). The main result is then the following.

Theorem 3.14. *If $h : \Omega \longrightarrow \Omega'$ is a homeomorphism between two tiling spaces of aperiodic and repetitive tilings with finite local complexity, then the associated complexity functions p and p' are equivalent in the sense above.*

Corollary. *Let T be an aperiodic, repetitive tiling with finite local complexity. Assume that its complexity function grows asymptotically like r^α. Then the exponent α is a topological invariant of its tiling space Ω_T.*

For subshifts, finite local complexity is automatic. Flow-equivalence of subshifts is defined as homeomorphism of their suspension. The following result is then an immediate consequence of the main theorem.

Corollary. *Two repetitive and aperiodic subshifts of $\mathbb{Z}^d$ which are flow-equivalent have equivalent complexity functions.*

From these results, it appears that complexity of a tiling is strongly related with the topology of the tiling space. Such links had already been investigated, to some extent. For example, for subshifts, it is known that a topological conjugacy is a so-called “sliding block code” (which is referred to as “local derivation” in our context, see Definition 2.3), and asymptotic properties of the complexity function are unchanged by an invertible recoding. For continuous tiling spaces, there is no longer an equivalence between topological conjugacy and local derivation. Frank and Sadun [16] proved in a more general setting (including even infinite local complexity), that the asymptotic behaviour of the complexity function is preserved by a topological conjugacy of the tiling spaces. These are strong arguments for stating that the complexity function is a feature of the dynamical system.

When it comes to how complexity of a tiling relates to the topology of its space (leaving the dynamics aside), only qualitative links have been established. It was proved in [8] that for an aperiodic tiling of dimension 1, the cohomology groups of the tiling space are finitely generated over $\mathbb{Q}$ if the complexity function is bounded above by a linear function. If one restricts to cut-and-project tilings, the result is more general: an aperiodic cut-and-project tiling of dimension d has finitely generated cohomology groups over $\mathbb{Q}$ if and only if its complexity function is bounded above by Cr^d for some constant C.

The methods for proving Theorem 3.14 (namely Lemma 3.9) are somehow related with results proved in [7, Chapitre 6]. It was proved that given a small deformation of a tiling, the deformed tiling and the original tiling have equivalent complexity functions. This result applied to a family of deformations introduced by Sadun and Williams. In [18], they proved that a tiling space is homeomorphic to the suspension of a subshift; in [7], it was proved that this homeomorphism preserves the complexity.

It turns out not to be a coincidence that methods which applied for the study of complexity of deformed tilings can be used here. In a last section, the links between homeomorphism and deformations are outlined. Given $h : \Omega \to \Omega'$, it defines an element in the mixed cohomology group $\mathrm{H}^1_{\mathrm{m}}(\Omega, \mathbb{R}^d)$ defined by Kellendonk [10]. This group describes deformations up to topological conjugacy.

Acknowledgements

I wish to thank Ian Putnam, Johannes Kellendonk, Lorenzo Sadun and Christian Skau for useful discussions. I also wish to thank Michael Baake who gave me the opportunity to present an early version of this work at the University of Bielefeld. Part of the early work on this topic was realized during a previous post-doc at the University of Victoria, under the supervision of Ian Putnam, partially funded by the Pacific Institute for the Mathematical Sciences.

2 Preliminaries

2.1 Tilings and tiling spaces

In our discussion, a *tile* is a closed compact set of the Euclidean space $\mathbb{R}^d$, which is homeomorphic to a closed ball, and is the closure of its interior. A *prototile* is the equivalence class under translation of a tile. Note that a tile may be decorated (one can think of blue and red squares, for example). Formally, it would be a pair (t, l), where t is the tile itself (a subset of $\mathbb{R}^d$) and l is a label. This notation will remain implicit, and a tile will just be denoted by t.

Let $\mathcal{A}$ be a finite collection of prototiles, which is fixed from this point on. A tiling with prototiles in $\mathcal{A}$ is a collection of tiles $T = \{t_i\}_{i \in I}$, such that for all i, the translation class of t_i belongs to $\mathcal{A}$. Furthermore, the union of the tiles is all of $\mathbb{R}^d$, and the tiles may only intersect on their boundaries: $t_i \cap t_j$ has no interior, unless $i = j$.

Given a tile $t \subset \mathbb{R}^d$, and $x \in \mathbb{R}^d$, the tile $t - x$ is defined as the translate of t by x. In a straightforward way, $T - x$ is defined as $\{t_i - x\}_{i \in I}$. It is a tiling, and is in general different from T (in our discussion, tilings are *not* being identified when they are translates of each other).

A tiling T is *periodic* if there exists $x \in \mathbb{R}^d \setminus \{0\}$, such that $T - x = T$. It is called *non-periodic* otherwise.

Given a non-periodic tiling, it makes sense to study the set of all tilings which cannot be distinguished from it at the local scale. This is the motivation for constructing the tiling space, in the same way one can build a subshift associated with a word.

Definition 2.1. Given T a tiling of $\mathbb{R}^d$, made from a set of prototiles $\mathcal{A}$, the *hull* of T is the following closure:

$$\Omega_T = \overline{\{T - x \, ; \, x \in \mathbb{R}^d\}}.$$

The closure is taken in the set Ω_{all} of all tilings made from tiles in $\mathcal{A}$, for the topology of the distance defined by:

$$\begin{aligned} D(T_1, T_2) \leq \varepsilon \text{ if } \exists x_1, x_2 \in \mathbb{R}^d, & \text{ with } \|x_i\| \leq \varepsilon, \\ & \text{such that } (T_1 - x_1) \text{ agrees with } (T_2 - x_2) \text{ on } B(0, 1/\varepsilon), \end{aligned}$$

and D is bounded above by $1/\sqrt{2}$.

Note that all elements in Ω_T are tilings themselves, and therefore there is an action of $\mathbb{R}^d$ by translation on Ω_T. This action is continuous for the topology defined above.

A tiling T is said to have *finite local complexity* (or FLC) if for all $R > 0$ there is a finite number (up to translation) of patches of the form $T \cap B(x, R)$, for $x \in \mathbb{R}^d$. Here, $T \cap B(x, R)$ stands for the patch of all tiles which intersect the ball $B(x, R)$.

A tiling T is said to be *repetitive* if for all finite patch $P \subset T$, there exists $R > 0$ such that a translate of P appears in $T \cap B(x, R)$, for any $x \in \mathbb{R}^d$.

The following proposition is well known.

Proposition 2.2. *Let T be a non-periodic and repetitive tiling, with finite local complexity. Then the tiling space Ω_T is a compact space, and the action by translation is minimal. When there is no risk of ambiguity, the tiling space is just written Ω.*

Note that all these definitions, which hold for tilings, would also hold for other patterns, as long as it is possible to define what a patch is, and what "identical up to translation" means. Delaunay sets are an important example. A Delaunay set of $\mathbb{R}^d$ is a subset $X \subset \mathbb{R}^d$, which is *uniformly discrete* and *relatively dense*. These two conditions mean respectively that there are two constants m and M such that any two points in X are separated by a distance at least m; and any point in $\mathbb{R}^d$ is within distance at most M of a point in X. The definition of *patch* (or local configuration) of X now becomes transparent when using the notation $X \cap B(x, r)$. From there, the definitions of FLC, repetitivity and aperiodicity, as well as the definition of the hull (and its properties) are unchanged.

2.2 Local derivations, and transversals

Given a tiling T, the associated tiling space Ω_T is sometimes called the "continuous" hull. It is also possible to describe a "discrete" tiling space, which is the analogue of a subshift in the symbolic setting (see Section 2.4). Amongst other properties, it is a subset of Ω, it intersects all $\mathbb{R}^d$-orbit and its intersection with any orbit is countable. It does not carry a $\mathbb{Z}^d$-action in general. This discrete tiling space is sometimes called a *canonical transversal*. The canonicity of this transversal leaves however a lot of room for choice.

First, define what a local derivation is.

Definition 2.3 (see [2])**.** Let Ω and Ω' be tiling spaces of repetitive and FLC tilings. A *local derivation* is a factor map (*i.e.* it is onto and commutes with translations) $\varphi : \Omega \to \Omega'$ which satisfies the following local condition:

$$\exists r, R > 0,\ \forall T \in \Omega,\ \varphi(T) \cap B(0, r) \text{ only depends on } T \cap B(0, R).$$

These conditions really mean that φ is a map defined on patches of size R, and $\varphi(T)$ is obtained by gluing together the images under φ of the patches out of which T is made. In this case, $\varphi(T)$ is said to be locally derived from T. When φ is invertible and φ^{-1} is also a local derivation, T and $\varphi(T)$ are called *mutually locally derivable*, or MLD.

Definition 2.4. Let Ω be a tiling space, and $\mathcal{D}$ be a local derivation rule defined on Ω, such that for all $T \in \Omega$, $\mathcal{D}(T)$ is a Delaunay set. Such a rule is called a *local pointing rule*. The *canonical transversal* associated with $\mathcal{D}$ is the set

$$\Xi_{\mathcal{D}} = \{T \in \Omega\,;\, 0 \in \mathcal{D}(T)\}.$$

The Delaunay set $\mathcal{D}(T)$ is repetitive and has finite local complexity, since T has these properties itself. It does not need to be MLD to T, or even aperiodic (see Figure 1).

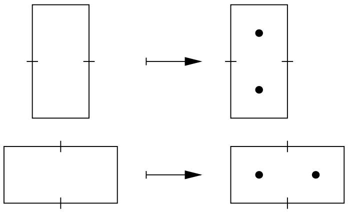

Figure 1: A local pointing rule, as defined on a set of rectangular tiles of size 2 × 1. Given any domino tiling made from these tiles (with dominoes meeting full-edge to full-edge), its associated point pattern will be periodic, even when the tiling is not.

Usually, what is called "canonical transversal" is built by choosing a pointing rule $\mathcal{D}$ which selects exactly one point in the interior of each prototile (for example their barycentre if the tiles are convex), and by pointing the tilings consistently.

Canonical transversals in the sense above are abstract transversals in the sense of Muhly Renault and Williams [15]. The reason why they are called "canonical" in spite of their apparent lack of canonicity is because T and T' are close in a canonical transversal if they agree *exactly* on a large ball centred at the origin (not up to a small translation).

Remark 2.5. If Ξ is a canonical transversal in a tiling space Ω, the topology on Ξ is induced by the family of clopen sets:

$$U_P = \{T \in \Xi\,;\, P \subset T\},$$

where P is in the family of admissible patches. Equivalently, it is induced by the metric:

$$D'(T,T') = \inf\Big(\{\varepsilon > 0\,;\, T \cap B(0,\varepsilon^{-1}) = T' \cap B(0,\varepsilon^{-1})\} \cup \{1/\sqrt{2}\}\Big).$$

Another description of the "canonical" transversals is that they are vertical with respect to a (truly) canonical solenoid structure on the tiling space.

Theorem 2.6 (see for ex. [3, 4])**.** *The tiling space associated with a repetitive, aperiodic and FLC tiling is an abstract solenoid in the following sense: for each* $T \in \Omega$*, there is a neighbourhood* V *of* T*, and a chart map* φ*, which maps* V *homeomorphically to the direct product of a Cantor set by a ball of* $\mathbb{R}^d$*. Furthermore,* φ *commutes with the action of* $\mathbb{R}^d$ *by translation whenever this action is defined. Also, the transition maps satisfy the property:*

$$\psi \circ \varphi^{-1}(\xi, x) = (\xi', x - t_{\psi\circ\varphi^{-1}}), \tag{1}$$

where ξ' *depends on* ξ*, but the translation vector* $t_{\psi\circ\varphi^{-1}}$ *only depends on the chart maps.*

Proof. This is a partial proof in order to outline how the canonical solenoid structure is chosen. Property (1) is not be proved here.

Let $T \in \Omega$. Let $P = T \cap B(0,1)$ be a small patch of T around the origin, and let $\varepsilon > 0$ be small with respect with the inner radius of the tiles. Define the derivation rule

$$\mathcal{D}(T') = \{x \in \mathbb{R}^d\,;\, P \subset T' - x\}.$$

It is clearly a local rule, and the associated transversal $\Xi_{\mathcal{D}}$ is the set of all tilings which coincide with T up to radius at least 1. If T is repetitive, aperiodic, and has finite local complexity, these are classic facts that $\Xi_{\mathcal{D}}$ is a Cantor set, and that $\mathcal{D}(T')$ is a Delaunay set for all T'.

Now, consider the map

$$\begin{array}{rccc} \varphi: & \Xi_{\mathcal{D}} \times \overline{B(0,\varepsilon)} & \longrightarrow & \Omega \\ & (T',x) & \longmapsto & T' - x. \end{array}$$

It is clearly continuous, by continuity of the action. It is one-to-one, because ε was chosen small. It is therefore a homeomorphism onto its image, by compactness of the source space. Finally, it is an exercise to check that the image of φ restricted to $\Xi_{\mathcal{D}} \times B(0,\varepsilon)$ is open in Ω. It provides the desired chart. It is important to remark that, by construction, the chart maps commute with the action of $\mathbb{R}^d$ whenever it is defined. □

It can be proved that the canonical transversals in the sense of Definition 2.4 are exactly the vertical transversals for this solenoid structure, that is: subsets $\Xi \subset \Omega$ which can be covered by chart boxes, and such that in each box $V \simeq X_V \times B(0,\varepsilon_V)$,

$$\Xi \cap V \simeq X \times \{0\}.$$

2.3 Complexity

The complexity function of a tiling T is a function p, such that $p(r)$ counts the number of patches of size r in T. This require to define what “patch of size r” means in this context. While, in the 1-dimensional symbolic setting, there is no question about what a word of length n is, the situation is not so obvious in the continuous case. However, it appears that the asymptotic properties of the complexity function are unchanged by picking any reasonable definition of “patch of size r”.

First, remark that a tiling T needs to have FLC in order to have a well defined complexity function (at least, in terms of naively counting patches—see [16] for ways of defining complexity for tilings without FLC). Second, if a tiling T is repetitive, then all tilings in its hull have the same complexity function. Therefore, for minimal tiling spaces, a complexity function is associated with the space rather than with one specific tiling. This justifies the following definition.

Definition 2.7. Let Ω be a space of tilings which are repetitive, FLC and aperiodic. Let $\mathcal{D}$ be a local pointing rule, and Ξ the associated transversal.

Then, the complexity function associated with these data, noted p_Ξ or $p_{\mathcal{D}}$ is:

$$\begin{aligned} p_\Xi(r) &= \mathrm{Card}\{T \cap B(0,r)\,;\, T \in \Xi\} \\ &= \mathrm{Card}\{(T_0 - x) \cap B(0,r)\,;\, x \in \mathcal{D}(T_0)\} \text{ for any } T_0 \in \Omega. \end{aligned}$$

According to this definition, the complexity function counts the number of *pointed patches* of size n, where the pointing has to be chosen in advance, and is not canonical to the tiling. However, it turns out that changing the transversal doesn't change the complexity function.

Proposition 2.8. *Let Ω be a tiling space, and Ξ_1, Ξ_2 be two canonical transversals associated with pointing rules $\mathcal{D}_1$ and $\mathcal{D}_2$ respectively. Then the complexity functions p_{Ξ_1} and p_{Ξ_2} are equivalent in the following sense:*

$$p_{\Xi_2}(r-c) \leq p_{\Xi_1}(r) \leq p_{\Xi_2}(r+C)$$

for some constants c, C and all r big enough.

Proof. Let $T \in \Omega$. First, let $\mathcal{D}$ be the pointing rule defined by

$$\mathcal{D}(T) = \mathcal{D}_1(T) \cup \mathcal{D}_2(T).$$

It is a pointing rule itself: $\mathcal{D}(T)$ is locally derived from T, it is relatively dense, and the uniform discreteness comes from finite local complexity of T (and hence, finite local complexity of $\mathcal{D}(T)$).

Let Ξ be the transversal defined by $\mathcal{D}$. It is sufficient for proving the theorem to show that p_Ξ is equivalent (in the sense above) to p_{Ξ_i}, for $i = 1, 2$. Let us prove it for $i = 1$.

For any given r, let X_r be the set of all patches of the form $(T-x) \cap B(0,r)$, with $x \in \mathcal{D}(T)$, and Y_r be the similarly defined set with $x \in \mathcal{D}_1(T)$. First, it is obvious that $p_{\Xi_1}(r) \leq p_\Xi(r)$ for all r (since $Y_r \subset X_r$). Conversely, let R be the constant such that for all $x \in \mathbb{R}^d$, $B(x,R)$ intersects $\mathcal{D}_1(T)$. Then, for all patch $P \in X_r$, there is a patch $P' \in Y_{r+R}$ such that P is a subpatch of P' (more precisely, there is $x \in B(0,R)$ such that $P + x \subset P'$, see Figure 2). Therefore,

$$p_\Xi(r) \leq p_{\Xi_1}(r+C).$$

□

Since the choice of a transversal does not change the asymptotic behaviour of the complexity function, it is acceptable to write "the complexity function of T" when there is no risk of ambiguity. It is understood that an unspecified transversal is chosen, the choice of which is not important.

Remark 2.9. It can also be proved that these pointed complexity function p_Ξ are equivalent to an unpointed complexity function p, where $p(r)$ is the number of equivalence classes up to translation of patches of the form $T \cap B(0,r)$, for $T \in \Omega$. It is proved in [7, Proposition 1.1.20].

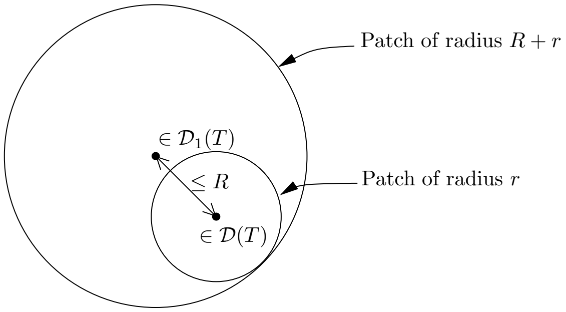

Figure 2: $\mathcal{D}_1$ is sparser than $\mathcal{D}$; yet, any patch of size r centred at a point of $\mathcal{D}$ is a subpatch of some patch of size $R + r$ centred at a point of $\mathcal{D}_1$.

Changing the transversal does not change the complexity. It is worth noticing that applying an invertible local derivation to a tiling also preserves the complexity function. It can be proved directly by remarking that a local derivation is essentially a map from patches to patches. The proof is not included here. See Theorem 3.7 for a more general result.

Proposition 2.10. *Let $T \in \Omega$ and $T' \in \Omega'$ be two tilings which are mutually locally derivable. Then their complexity functions are equivalent:*

$$p(r - c) \leq p'(r) \leq p(r + C),$$

for r big enough.

The notions of equivalence which are used in the propositions above are quite fine. It is appropriate for our purpose to have a coarser notion of equivalence between complexity functions.

Definition 2.11. Two complexity functions p_1 and p_2 are *equivalent* if there are $C_1, C_2, m, M > 0$ such that for all r big enough:

$$C_1 p_2(mr) \leq p_1(r) \leq C_2 p_2(Mr).$$

Remark that it is indeed a coarser notion that the ones above. However, the notion of a "patch of size r" is itself not uniquely defined: it depends on a choice of norm of $\mathbb{R}^d$. By equivalence of norms in finite dimension, two different choices of norms would produce equivalent complexity functions in the sense above. Therefore, and unless there is a reason to prefer one norm over others, it is fine to use this coarser noton of equivalence.

2.4 Formalism of symbolic dynamics

Tiling spaces also appear as suspension of subshifts. A d-dimensional subshift is a subset of $\mathcal{A}^{\mathbb{Z}^d}$, where $\mathcal{A}$ is a finite set of symbols (finite alphabet). There is a natural action σ of $\mathbb{Z}^d$ on this set, defined by

$$[\sigma^n(\omega)]_k = \omega_{k-n},$$

for $n \in \mathbb{Z}^d$.

Given $\omega \in \mathcal{A}^{\mathbb{Z}^d}$, the subshift associated with ω is the smaller subset of $\mathcal{A}^{\mathbb{Z}^d}$ which contains ω, is closed and is shift-invariant. Equivalently, it is the closure of the orbit of ω for the product topology.

The notions of aperiodicity and repetitivity are straightforward. Finite local complexity is automatic if $\mathcal{A}$ is finite.

Definition 2.12. Let Ξ be an aperiodic, repetitive subshift on $\mathbb{Z}^d$. The *suspension* of Ξ is the following space:

$$\mathcal{S}\Xi = \Xi \times \mathbb{R}^d / \sim$$

where $\sim$ is the equivalence relation generated by

$$(\omega, x) \sim (\sigma^n(\omega), x - n), \quad n \in \mathbb{Z}^d.$$

It carries naturally a $\mathbb{R}^d$-action by translations.

In practice, it is convenient to view elements of Ξ as tilings by d-cubes of edge-length 1, and $(\omega, x) \in \mathcal{S}\Xi$ can be interpreted as $\omega + x$. The tiling topology, as defined above, coincides with the quotient topology on this suspension.

Definition 2.13. Two subshifts Ξ_1 and Ξ_2 are said to be *flow-equivalent* if their suspensions are homeomorphic.

3 Homeomorphisms between tiling spaces

In this section, the consequence on the complexity function of a homeomorphism between two aperiodic, FLC, minimal tiling spaces is investigated.

3.1 Consequences on the Groupoid

Given a dynamical system, it is possible to define a groupoid which carries information on the space, the group, and the action of the latter on the former. Groupoids also serve as a generalization of group action, when there is no group. For example, a tiling space carries an action of $\mathbb{R}^d$. In one dimension, this action of $\mathbb{R}$ restricts to an action of $\mathbb{Z}$ on a transversal Ξ (by the first return map). In higher dimension, there is in general no group action on the transversal; however, a groupoid can be defined as a replacement.

Definition 3.1. A groupoid G is defined as a small category with inverses. It consists of a set of base points $G^{(0)}$, and a set of arrows G between them. Given $x \in G$, it has a source and a range $s(x)$ and $r(x)$ in $G^{(0)}$ and the product xy is well defined if $r(y) = s(x)$. For all $p \in G^{(0)}$ there is an element e_p which is neutral for multiplication (and p is identified with e_p, so that $G^{(0)} \subset G$), and for all x there is x^{-1} such that $xx^{-1} = r(x)$, $x^{-1}x = s(x)$. Finally, the product is associative, whenever it is defined.

Definition 3.2. A tiling space Ω has an associated groupoid, noted $\Omega \rtimes \mathbb{R}^d$, which is defined as the direct product of Ω and $\mathbb{R}^d$, with range map, source map and product defined as follows:

$$\begin{array}{cc} s(T,x) = T, & r(T,x) = T - x, \\ \multicolumn{2}{c}{(T,x)\cdot(T',x') = (T', x+x') \quad \text{if} \quad T' - x' = T.} \end{array}$$

The topology on this groupoid is that induced by $\Omega \times \mathbb{R}^d$.

Consider a homeomorphism $h : \Omega \longrightarrow \Omega'$ between two aperiodic, FLC and minimal tiling spaces. By Theorem 2.6, the path-connected components in such tiling spaces are the orbits. Since a homeomorphism has to send path-connected component to path-connected component, we can define the following map.

Definition 3.3. Let $h : \Omega \longrightarrow \Omega'$ be a homeomorphism. Then for all $T \in \Omega$, let $h_T : \mathbb{R}^d \to \mathbb{R}^d$ be the function defined by

$$h(T - x) = h(T) - h_T(x).$$

Theorem 3.4. *Given a homeomorphism $h : \Omega \longrightarrow \Omega'$, the function*

$$\begin{array}{ccc} \Omega \times \mathbb{R}^d & \longrightarrow & \mathbb{R}^d \\ (T,x) & \longmapsto & h_T(x) \end{array}$$

is continuous in two variables. Furthermore, for any fixed T, h_T is a homeomorphism.

Proof. The first step of this proof is to show that for any given $T \in \Omega$, the map $x \mapsto h_T(x)$ is continuous. The map h_T is defined as the map which makes the following diagram commute:

$$\begin{tikzcd} \mathbb{R}^d \arrow[r, "h_T"] \arrow[d] \arrow[dr, phantom, "\circlearrowleft"] & \mathbb{R}^d \arrow[d] \\ \mathrm{Orb}(T) \arrow[r, "h"] & \mathrm{Orb}(h(T)) \end{tikzcd}$$

where $\mathrm{Orb}(T)$ is the orbit of T in Ω under the action of $\mathbb{R}^d$. The difficult point is that the action maps $\mathbb{R}^d$ continuously and bijectively onto $\mathrm{Orb}(T)$ via $x \mapsto T - x$; however, for the relative topology induced by that of Ω, $\mathrm{Orb}(T)$ is not homeomorphic to $\mathbb{R}^d$, and not even locally homeomorphic. Indeed, the image of an open set is not open (in general) for the induced topology, and $\mathrm{Orb}(T)$ is not locally connected. However, Theorem 2.6 gives information about the local structure of the leaves, which allows to prove that h_T is a homeomorphism.

Here, the sequential characterization of continuity is used: T and x being fixed, let $(x_n)_{n\in\mathbb{N}}$ be a sequence converging to x. Let $\gamma : [0,1] \to \mathbb{R}^d$ be the continuous path from x_1 to x, such that γ restricted to $[1 - 2^{-n}, 1 - 2^{-(n+1)}]$ is the constant speed parametrization of the segment $[x_{n+1}, x_{n+2}]$. To prove continuity of h_T at x, it is enough to prove that $h_T \circ \gamma$ tends to $h_T(x)$ as t tends to 1.

Since h, γ and of the action of $\mathbb{R}^d$ on Ω are continuous, $h(T-\gamma(t))$ is close to $h(T-x)$ provided t is close to 1. Assume t_0 is close enough to 1 so that for all $t > t_0$, $h(T-\gamma(t))$ is in a chart box of the form $V \simeq_\varphi X \times B(0,\varepsilon)$ around $h(T-x)$, and such that $\varphi(h(T-x)) = (\xi, 0)$.

For $t > t_0$, the path $t \mapsto \varphi \circ h(T-\gamma(t))$ is a continuous path in $X \times B(0,\varepsilon)$, and therefore must be included in a path-connected component of the form $\{\xi'\} \times B(0,\varepsilon)$, and therefore be of the form $(\xi', \eta(t))$. Since $\varphi \circ h(T-\gamma(t))$ needs to tend to $(\xi, 0)$, it is actually the case that $\xi' = \xi$.

Finally,

$$\begin{aligned} \varphi \circ h(T-\gamma(t)) &= \varphi(h(T) - h_T \circ \gamma(t)) \subset \{\xi\} \times B(0,\varepsilon), \\ &= (\xi, \eta(t)). \end{aligned}$$

Now, $\eta(t)$ tends to zero on the one hand, and is equal to $h_T(\gamma(t)) - h_T(x)$ on the other hand (because φ commutes with the action of $\mathbb{R}^d$ whenever it is defined, see Theorem 2.6). Therefore, $h_T(\gamma(t))$ tends to $h_T(x)$, and h_T is continuous.

The second part of the proof is to show that the function is continuous of two variables. Let $(T_0, x_0) \in \Omega \times \mathbb{R}^d$. What needs to be shown is

$$\begin{aligned} &\forall \varepsilon > 0 \ ; \ \exists \delta > 0 \ ; \ \forall (T,x) \in \Omega \times \mathbb{R}^d, \\ &\qquad \Big(\|x - x_0\| < \delta \text{ and } D(T,T_0) < \delta \Big) \Rightarrow \|h_T(x) - h_{T_0}(x_0)\| < \varepsilon. \end{aligned}$$

Let $\varepsilon > 0$. Let δ_0 be such that whenever T_1, T_2 are δ_0-close, then their images by h are $\varepsilon/3$-close (using uniform continuity of h). Finally, let

$$\delta = \min \big\{ (\delta_0^{-1} + \|x_0\| + 2\delta_0)^{-1}, \delta_0/2 \big\}.$$

Let now $T \in \Omega$ in a δ-neighbourhood of T_0. Then, up two small translation of combined size less than $2\delta < \delta_0$, T and T_0 agree on $B(0, \delta_0^{-1} + \|x_0\| + \delta_0)$. Then, for all $y_0 \in \overline{B(0, \|x\|)}$ and all $y \in \mathbb{R}^d$ with $\|y - y_0\| < \delta \leq \delta_0/2$, the tilings $T - y$ and $T - y_0$ agree at least on $B(0, \delta_0)$ up to a small (less than δ) translation. This last statement holds in particular for $y = \lambda x$ ($\lambda \in [0,1]$).

Let $x \in B(x_0, \delta)$, and $T \in \Omega$ be δ-close from T_0. Then, for all $\lambda \in [0,1]$, $h(T - \lambda x)$ and $h(T_0 - \lambda x_0)$ are $\varepsilon/3$-close. Assume (by contradiction), that $\|h_T(x) - h_{T_0}(x_0)\| \geq \varepsilon$. Since h_T and h_{T_0} are continuous, and $\|h_T(0) - h_{T_0}(0)\| = 0$, one can use the intermediate value theorem: there exists a $\lambda_0 \in [0,1]$ such that $\|h_T(\lambda_0 x) - h_{T_0}(\lambda_0 x_0)\| = \varepsilon$. One has

$$\begin{aligned} h(T - \lambda_0 x) &= h(T) - h_T(\lambda_0 x) \\ h(T_0 - \lambda_0 x_0) &= h(T_0) - h_{T_0}(\lambda_0 x_0). \end{aligned}$$

On the one hand, $h(T)$ and $h(T_0)$ match around the origin up to a translation smaller than $\varepsilon/3$. On the other hand, so do the left-hand sides of the equations above. Besides, $h_T(\lambda_0 x)$ and $h_{T_0}(\lambda_0 x_0)$ differ by exactly ε. Since ε much smaller than the radius of a tile, it is a contradiction. □

Remark 3.5. The map $(T, x) \mapsto h_T(x)$ is actually a groupoid morphism (where the additive group $\mathbb{R}^d$ is seen as a groupoid with one basepoint). It results from a direct check:

$$h_T(x+y) = h_T(x) + h_{T-x}(y).$$

This relation can also be called a cocycle condition for a reason which will be addressed in the last section.

Remark 3.6. A consequence of this result is that the following map

$$\Omega \rtimes \mathbb{R}^d \longrightarrow \Omega' \rtimes \mathbb{R}^d (T, x) \longmapsto (h(T), h_T(x))$$

is an isomorphism of continuous groupoids. Therefore, by [17], the associated cross product C^*-algebras are isomorphic. An interpretation of this fact is that the local product structure of FLC tiling spaces forces strong links between topology and dynamics: if the topological spaces are homeomorphic, it implies that the groupoids and associated C^*-algebra—which encode the dynamics—are isomorphic.

3.2 Consequences for the complexity function

In this section, it proved that whenever there is a homeomorphism between two tiling spaces, the complexity functions are equivalent in the sense of Definition 2.11. If the complexity is high, then the tiling systems may have entropy. In this case, the entropy is not preserved by homeomorphism (even though the fact that it is positive is). However, if the complexity grows polynomially, then the exponent is preserved by homeomorphism.

Frank and Sadun have proved in [16] that whenever two tiling spaces are topologically conjugate, their complexity functions are equivalent. Their result involve a different complexity function as the one which is used here, as it can measure complexity of tiling spaces with infinite local complexity. Here is how their result would read in our framework.

Theorem 3.7 (Frank–Sadun). *Let $h : \Omega \longrightarrow \Omega'$ be a topological conjugacy between two tiling spaces (that is a homeomorphism which commutes with the actions of $\mathbb{R}^d$ by translation). We assume that the tiling spaces are aperiodic, minimal, with FLC. Then, there are two canonical transversals $\Xi \subset \Omega$ and $\Xi' \subset \Omega'$, such that the associated complexity function p and p' satisfy*

$$p(r-a) \leq p'(r) \leq p(r+b).$$

for some constants a and b, and all r big enough.

Note that by Proposition 2.8, it is not important which transversals Ξ and Ξ' are chosen, as any complexity functions would then satisfy this inequality.

We include a sketch of the proof, with a simplifying assumption that h sends a vertical transversal to a vertical transversal. Even in this simplified setting, the proof contains the main ideas which will be used in the proof of Theorem 3.14. It will be pointed out how they should be generalized.

Proof. Let $\Xi \subset \Omega$ be any canonical transversal. It is assumed that $\Xi' := h(\Xi)$ is a canonical transversal itself (this is a simplifying assumption, as it shouldn't be expected in general).

The goal is then to compare the functions p and p' associated with Ξ and Ξ' respectively. Let $\varepsilon < 2^{-1/2}$. By uniform continuity of h, there is a $\delta < 2^{-1/2}$ such that for all $T, T' \in \Xi$, if T and T' are at least δ-close, then $h(T)$ and $h(T')$ are at least ε-close. In terms of patches, it means that if T and T' agree on $B(0, \delta^{-1})$, then their images agree on $B(0, \varepsilon^{-1})$. In particular, $h(T)$ and $h(T')$ have the same local configuration at the origin.

Now, ε and δ being fixed, let $r > 0$, big enough so that it is big compared to ε^{-1} and δ^{-1}. Assume T and T' agree on a ball of radius $r + \varepsilon^{-1}$. Then it means that for all $x \in \mathbb{R}^d$ with $\|x\| \leq r$, $T + x$ and $T' + x$ agree on $B(0, \delta^{-1})$. Using the fact that h is a conjugacy, it means that for all such x, $h(T) + x$ and $h(T') + x$ agree on $B(0, \varepsilon^{-1})$. Therefore, T and T' agree up to radius $r + \delta^{-1}$ implies that their images agree up to radius $r + \varepsilon^{-1}$. Therefore,

$$p(r + \delta^{-1}) \geq p'(r + \varepsilon^{-1}).$$

Exchanging h and h^{-1}, one gets the other inequality. □

This proof has three essential ingredients:

1. The image of a transversal is a transversal (simplifying assumption);
2. The map h is uniformly continuous, and it allows to control how two tilings agree on an immediate neighbourhood of the origin when their preimages do;
3. The conjugacy sends orbits to orbits in a trivial way (given T, the map h_T is the identity map), so that patches of tilings in Ω can be compared to patches of tilings in Ω'.

In the case when h is just a homeomorphism, point (2) is still satisfied. Point (1) is not satisfied, however it is very close to being satisfied: for all $\varepsilon > 0$, there are transversals $\Xi \subset \Omega$ and $\Xi' \subset \Omega'$ such that $h(\Xi)$ and Ξ' are within distance ε.

The main difficulty is point (3). Given a ball of radius r in an orbit, what is the radius of its image by h? We need to control how $h_T(B(0, r))$ grows with r. Lemma 3.9 and its corollary shows that h_T is not far from being Lipschitz.

The following lemma addresses point (1) above: a homeomorphism between two tiling spaces is very close to sending canonical transversal to canonical transversal.

Lemma 3.8. Let $h : \Omega \longrightarrow \Omega'$ be a homeomorphism. Then for all $\varepsilon > 0$, there is a vertical transversal $\Xi \subset \Omega$, a vertical transversal $\Xi' \subset \Omega'$, and a homeomorphism $h_\varepsilon : \Omega \longrightarrow \Omega'$ such that

- $h_\varepsilon(\Xi) = \Xi'$;
- h and h_ε are isotopic;

- there is a continuous function $s : \Omega \to \mathbb{R}^d$, which is bounded by ε, such that $h_\varepsilon = h - s$.

Proof. Let $\varepsilon > 0$, which is assumed to be small compared to the inner radius of the tiles. By uniform continuity of h, there is a $\delta > 0$ such that $d(T, T') \leq \delta$ implies $d(h(T), h(T')) \leq \varepsilon$. Then, there is a chart box $V_0 \subset \Omega'$, of diameter less than ε. Let U_0 be a chart box in Ω such that $U_0 \subset h^{-1}(V_0)$. Then we have the following picture for some $\delta' \leq \delta$ and $\varepsilon' \leq \varepsilon$.

$$\begin{tikzcd}
U_0 \arrow[r, "h"] \arrow[d, "\simeq"'] & V_0 \arrow[d, "\simeq"] \\
X_0 \times B(0, \delta') \arrow[r, "\varphi"] & X'_0 \times B(0, \varepsilon')
\end{tikzcd}$$

where φ is induced by h (so that the diagram commutes). See also Figure 3. When no confusion can arise, φ will not be used, and this map will just be noted h, possibly with the precision "in coordinates".

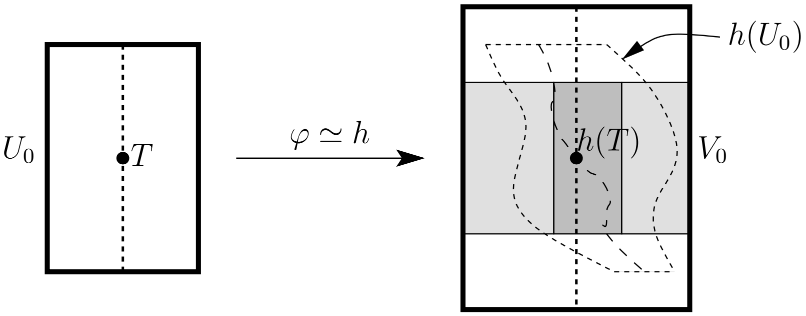

Figure 3: U_0 and V_0, as initially built. They need pruning: for example V_0 is too tall. First, pick a box-shaped neighbourhood of $h(T)$ inside $h(U_0)$ (dark grey neighbourhood). Then, widen it (light-grey neighbourhood). Finally, trim U_0 so that its image is now included in the light-grey neighbourhood.

With a good choice of charts, we can assume that $\varphi(\xi, 0) = (\xi', 0)$ for some $\xi \in X_0$ and $\xi' \in X'_0$. Call T_0 and T'_0 the tilings corresponding to $(\xi, 0)$ and $(\xi', 0)$ respectively via the chart maps (so that $h(T_0) = T'_0$). Note that $h(U_0)$ is open in V_0. Therefore, there is a neighbourhood of T'_0, $V' \simeq X' \times B(0, \varepsilon'_1)$ such that $V' \subset h(U_0)$. Let now $V \simeq X' \times B(0, \varepsilon')$, and U be the largest open subset of U_0 such that $\varphi(U) \subset V$. It is easy to check that it is of the form $X \times B(0, \delta')$, with X a clopen subset of X_0.

Let Ξ and Ξ' be the vertical transversals included in U and V respectively, such that $\Xi \simeq X \times \{0\}$ and $\Xi' \simeq X' \times \{0\}$.

We now have U and V two chart domains with $h(U) \subset V$, and two transversals $\Xi \subset U$ and $\Xi' \subset V$ such that $\Xi' \subset h(U)$.

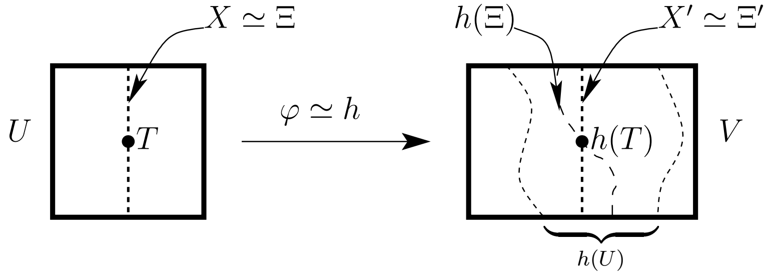

Figure 4: The construction of Ξ and Ξ'. The diameter of U is at most δ and the diameter of V is at most ε.

The next step in the proof is to build a map $\Xi \to \Xi'$ which is induced by h. The definition itself is straightforward: for each $T \in \Xi$, there is a unique $T' \in \Xi'$ which belongs to the same connected component of V as $h(T)$. Define $h_\varepsilon(T)$ to be this T'. It is clear that for all $T \in \Xi$, $h_\varepsilon(T) = h(T) - s(T)$, where $s(T)$ is a translation vector of norm less than ε (since the diameter of V is less than ε). Two things need to be shown: first, that h_ε is indeed continuous; then, that it can be extended to a function on all Ω and not just on Ξ.

To this end, it is proved that s is a continuous function. Let $\eta > 0$, which is assumed to be small with respect to the size of the tiles, and $T_1 \in \Xi$. By continuity of h, there is $\nu > 0$ such that for all $T \in \Xi$ which is ν-close to T_1, the images by h of T and T_1 are $\eta/2$-close. It implies that there is an $x \in \mathbb{R}^d$, with $\|x\| < \eta$, such that $h(T_0)$ and $h(T) - x$ match exactly around the origin. In addition, since $h_\varepsilon(T) = h(T) - s(T)$ and $h_\varepsilon(T_0) = h(T_0) - s(T_0)$ both belong to Ξ', they need to match exactly around the origin (because ε is small). Therefore, *when restricted to* $B(0, \varepsilon^{-1})$, one has (the notation $=_r$ stands for "agree around the origin up to radius r"):

$$\begin{aligned} h(T_0) - s(T_0) &=_{1/\varepsilon} h(T) - s(T) \quad \text{so} \\ h(T_0) - s(T_0) &=_{1/\varepsilon} h(T) - (x + s(T_0)), \end{aligned}$$

and since $\|x\|$ and $s(T_0)$ are both smaller than ε (which is small compared to the size of tiles), the only way that the last two tilings can match around the origin is if $s(T) = s(T_0) + x$, which implies $\|s(T_0) - s(T)\| < \eta$. Since η could have been chosen arbitrarily small, it proves that s is continuous on Ξ.

This last fact provides an instant proof of the continuity of h_ε on Ξ: indeed, $h_\varepsilon = h - s$, both h and s are continuous, and $\mathbb{R}^d$ acts continuously on Ω by translations.

Let us now show that h_ε can be extended on Ω. It is done by extending s. First, φ (induced by h) extends to a map defined on the closure of its domain (by continuity of h on Ω). It is still noted φ:

$$\varphi : X \times \overline{B(0, \delta')} \longrightarrow X' \times \overline{B(0, \varepsilon')}$$

Then, for all $\xi \in X$, there is a map $\varphi_\xi : x \mapsto \varphi(\xi, x)$ defined on the closed ball of radius δ', which is a homeomorphism onto its image. By construction, its image contains 0 in its interior. Let $x_\xi := \varphi_\xi^{-1}(0)$. It is an inner point of the closed disk. Then, let us explicitly provide a homeomorphism f_ξ from the closed disk to itself, which leaves the boundary invariant and sends 0 to x_ξ.

Let $a(y) = 1 - \|y\| / \delta'$ defined on $\overline{B}(0, \delta')$. Notice that it is 0 on the boundary, and takes value 1 at 0. Now, define

$$f_\xi(y) = y + a(y) x_\xi. \tag{2}$$

It is clear that $f_\xi(0) = x_\xi$, and f_ξ leaves the boundary invariant. Furthermore, $y \mapsto a(y)x_\xi$ is k-Lipschitz, with $k = \|x_\xi\| / \delta' < 1$. Therefore, f_ξ is one-to-one, and so it is a homeomorphism onto its image.

Now, one can check that the norm of $f_\xi(y)$ is

$$\begin{aligned} \|f_\xi(y)\| &= \|y\| + \|x_\xi\| - \|y\| \|x_\xi\| / \delta' \\ &= \delta' \left(\frac{\|y\|}{\delta'} + \frac{\|x_\xi\|}{\delta'} - \frac{\|y\|}{\delta'} \frac{\|x_\xi\|}{\delta'} \right) \end{aligned}$$

and a simple study of the function $(\alpha, \beta) \mapsto \alpha + \beta - \alpha\beta$ on the domain $[0,1]^2$ shows that it is bounded by one. Therefore, the range of f_ξ is included in the closed ball $\overline{B}(0, \delta')$, and it contains its boundary. By an argument of algebraic topology[1], f_ξ is therefore onto, and is a homeomorphism of the closed ball.

It is a quick check that the function $\xi \mapsto f_\xi$ is continuous on Ξ. Indeed, the continuity of h_ε implies that $(\xi, x_\xi) = h^{-1} \circ h_\varepsilon(\xi, 0)$ is a continuous function of ξ (expressed in coordinates, the chart maps being implicit). In addition, the map $x \mapsto (\mathrm{id}(\cdot) + a(\cdot)x)$ is continuous for the topology of uniform convergence on the space of functions on the ball. It gives continuity of $\xi \mapsto f_\xi$.

Now, it is possible to extend the function h_ε on $U \simeq X \times B(0, \delta')$: in coordinates, it is defined as $h_\varepsilon(\xi, y) = (\xi, f_\xi(y))$. It amounts to precomposing h with f_ξ on each leaf. It is clear by definition of the f_ξ that this map restricts to h_ε on Ξ, as expected. The continuity of $\xi \mapsto f_\xi$ shows that h_ε is continuous on U. Finally, because of the boundary conditions put on f_ξ, the function h_ε can be extended continuously by h on the complement of U.

The last point to prove is the isotopy between h_ε and h. It can be done by replacing the definition of f_ξ in Equation (2) by the following one:

$$f_{\xi,t}(y) = y + t a(y) x_\xi,$$

for a parameter $t \in [0,1]$. Leaving the rest of the construction unchanged, the result is a family of functions $(h_{t\varepsilon})_{t \in [0,1]}$. For $t = 1$, the function is just h_ε; for $t = 0$, it is $h_0 = h$. This provides the desired isotopy. $\square$

The function h_ε being a well defined homeomorphism $\Omega \to \Omega'$, it is possible for all $T \in \Omega$ to define $h_{\varepsilon,T} : \mathbb{R}^d \to \mathbb{R}^d$ in the same way h_T was defined.

---

[1] Assume w in the interior of the ball is not on the range of f_ξ. Then f_ξ composed with $z \mapsto \delta'(z - w) / \|x - w\|$ maps the closed ball onto the closed sphere continuously, which is a contradiction.

Lemma 3.9. There is a $M > 0$ such that for all $T \in \Xi$ and all $x \in \mathbb{R}^d$ satisfying $T - x \in \Xi$,

$$\|h_{\varepsilon,T}(x)\| \leq M \|x\| .$$

Proof. Let $T \in \Xi$ and $T' = h_\varepsilon(T) \in \Xi'$. Let $\mathcal{D}$ and $\mathcal{D}'$ be the pointing rules associated with transversal Ξ and Ξ' respectively. It immediate to check that $x \in \mathcal{D}(T)$ if and only if $h_{\varepsilon,T}(x) \in \mathcal{D}'(T')$.

Let $\mathcal{T}$ be a triangulation locally derived from $\mathcal{D}$, such that the set of vertices is exactly $\mathcal{D}$. A triangulation is a tiling where each tile is a d-simplex, and two tiles meet on a $d-1$-simplex. One well known way to do so is to use Delaunay triangulations of $\mathbb{R}^d$ with vertices in $\mathcal{D}(T)$. Such triangulations are in general dual from the Voronoi construction. In some (non-generic) cases, the Delaunay triangulation is not unique, but by making consistent choices, it is possible to ensure that $\mathcal{T}(T)$ is locally derived from T. See [20], for a reference.

Let $g : \mathbb{R}^d \to \mathbb{R}^d$ be the function defined as follows: given $x \in \mathbb{R}^d$, it belongs to a simplex, so it is uniquely written as a convex combination (with non-negative coefficients) of its vertices, which are elements in $\mathcal{D}$:

$$x = \lambda_0 v_0 + \lambda_1 v_1 + \ldots + \lambda_d v_d, \quad \lambda_k \geq 0, \ \sum_{k=0}^{d} \lambda_k = 1.$$

Then, define

$$g(x) = \sum_{k=0}^{d} \lambda_k h_{\varepsilon,T}(v_k).$$

This map g is onto, but does not need to be one-to-one, see Figure 5. It is affine by parts, and using finite local complexity of $\mathcal{T}$, there are finitely may linear maps underlying the affine maps. A consequence of this is that g is Lipschitz, with coefficient M. Finally, g agrees with $h_{\varepsilon,T}$ on $\mathcal{D}$.

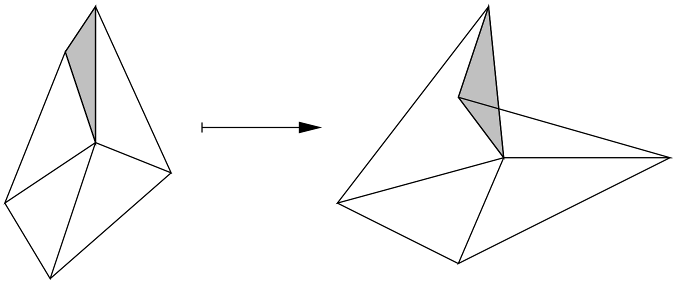

Figure 5: A local pattern of the triangulation can be sent to a pattern which is combinatorially a triangulation, but with an overlap in the geometric realization. Remark that the shaded triangle has a “wrong” orientation on the right. In this case, the map g is not one-to-one.

Now, let us prove the lemma. Let $x \in \mathbb{R}^d$ such that $T - x \in \Xi$ (that is: let $x \in \mathcal{D}$). Then

$$\begin{aligned} \|h_{\varepsilon,T}(x)\| &= \|h_{\varepsilon,T}(x) - h_{\varepsilon,T}(0)\| \\ &= \|g(x) - g(0)\| \\ &\leq M\|x - 0\| = M\|x\|. \end{aligned}$$

□

This lemma has the following consequence.

Corollary 3.10. There is a $M > 0$ and a $C > 0$ such that for all $T \in \Omega$ and all $x \in \mathbb{R}^d$,

$$\|h_{\varepsilon,T}(x)\| \leq M\|x\| + C.$$

Proof. Let Ξ and Ξ' be the vertical transversal involved in the construction of h_ε, and let $\mathcal{D}$ be the local pointing rule associated with Ξ. Using the Delaunay property of $\mathcal{D}(T)$ for $T \in \Omega$, there is $R > 0$ such that any ball $B(x, R)$ intersects $\mathcal{D}(T)$, for any $T \in \Omega$.

Using continuity of $(T, x) \mapsto h_{\varepsilon,T}(x)$ (see Theorem 3.4), the image of the compact set $\Omega \times \overline{B(0, R)}$ is compact, and in particular is bounded. It is therefore included in $B(0, c)$.

Let now be $x \in \mathbb{R}^d$ and $T \in \Omega$. Then there are $y, z \in \mathbb{R}^d$ of length less than R such that $T - y \in \Xi$ and $T - x - z \in \Xi$. One can write:

$$\begin{aligned} \|h_T(x)\| &= \|h_T(y + (x + z - y) - z)\| \\ &= \|h_T(y) + h_{T-y}(x + z - y) + h_{T-x-z}(-z)\| \\ &\leq c + M(\|x\| + \|y\| + \|z\|) + c \\ &\leq 2c + 2MR + M\|x\|. \end{aligned}$$

The equality from first line to second uses the definition of h_T, the first inequality follows from Lemma 3.9. The result follows, with $C = 2c + 2MR$. □

This previous lemma is the key for estimating the size of $h_T(B(0, R))$ as R grows. A few additional lemmas still need to be proved before tackling the main theorem.

The two following lemmas are treat the following problems: how can a given $x \in \mathbb{R}^d$ be approximated by a sum of small "return vectors" (Definition 3.12); and how $h_{\varepsilon,T}(v)$ is locally constant in T when v is a small return vector.

Lemma 3.11. Let $\mathcal{D}$ be a Delaunay set which is R-relatively dense, that is:

$$\bigcup_{p \in \mathcal{D}} B(p, R) = \mathbb{R}^d.$$

Let $x \in \mathcal{D}$ and $y \in \mathbb{R}^d$, such that $\|x - y\| \geq 2R$. Then there exists $x' \in \mathcal{D}$, such that

$$\|x' - x\| \leq 3R \quad \text{and} \quad \|x' - y\| \leq \|x - y\| - R.$$

Proof. Define

$$z = x - 2R\frac{x-y}{\|x-y\|}.$$

By the Delaunay property, there exists $x' \in \mathcal{D}$ which is within distance R of z. It is a quick check that $\|x - x'\| \leq 3R$, and

$$\begin{aligned} \|x' - y\| &\leq \|x' - z\| + \|z - y\| \\ &\leq R + \|x - y\| - 2R \\ &= \|x - y\| - R. \end{aligned}$$

□

Definition 3.12. Let Ω be a tiling space, and Ξ be a canonical transversal associated with the local pointing rule $\mathcal{D}$. Then the set

$$\mathcal{V} = \{v \in \mathbb{R}^d \,;\, \exists T \in \Xi,\ T - v \in \Xi\}$$

is discrete (by FLC), and is called the set of return vectors on Ξ. By definition,

$$\mathcal{V} = \bigcup_{T \in \Xi} \mathcal{D}(T).$$

Given $v \in \mathcal{V}$, define

$$\Xi_v = \{T \in \Xi \,;\, T - v \in \Xi\}.$$

By finite local complexity, it is a clopen subset of Ξ.

Lemma 3.13. Let $h_\varepsilon : \Omega \longrightarrow \Omega'$ be a homeomorphism, which sends a transversal Ξ to a transversal Ξ'. Let $\mathcal{V}$ be the set of return vectors on Ξ.

Let $\lambda > 0$, and $\mathcal{V}_\lambda$ be the set of all return vectors of length at most λ. Then there exists $\nu > 0$ such that for all $v \in \mathcal{V}_\lambda$ and all $T_1, T_2 \in \Xi_v$,

$$D(T_1, T_2) \leq \nu \Longrightarrow h_{\varepsilon,T_1}(v) = h_{\varepsilon,T_2}(v).$$

Proof. Given $v \in \mathcal{V}_\lambda$, the map

$$T \mapsto h_T(v)$$

is continuous on Ξ_v, and it maps v to return vectors on Ξ' (indeed, by definition of Ξ_v, both T and $T - v$ are in Ξ, and therefore their images are in Ξ'). It is a continuous map from a compact set to a discrete set (by finite local complexity). It has therefore finite range, and is locally constant. This gives a constant ν_v, such that

$$D(T_1, T_2) \leq \nu_v \Longrightarrow h_{\varepsilon,T_1}(v) = h_{\varepsilon,T_2}(v).$$

Since $\mathcal{V}_\lambda$ is finite, the lemma follows. □

All the tools are now available for proving the main theorem on complexity.

Theorem 3.14. *Let $h : \Omega \longrightarrow \Omega'$ be a homeomorphism between two aperiodic repetitive tiling spaces with finite local complexity. Then there exist a transversal $\Xi \subset \Omega$ and a transversal $\Xi' \subset \Omega'$ and two constants $m, M > 0$ such that for all r big enough,*

$$p_{\Xi'}(mr) \leq p_\Xi(r) \leq p_{\Xi'}(Mr).$$

Proof. Let Ξ, Ξ' and h_ε be given by Lemma 3.8. Denote by $\mathcal{D}$ and $\mathcal{D}'$ the local pointing rules associated with Ξ and Ξ'.

The strategy of this proof is to use Lemma 3.9 to prove that whenever $T, T' \in \Xi$ disagree within radius r, then their images in Ξ' need to disagree within radius at most Mr, for some constant M which does not depend on r. It will prove that

$$p_\Xi(r) \leq p_{\Xi'}(Mr).$$

The other inequality will be obtained by repeating the same argument for h_ε^{-1}.

Let R be the constant such that the Delaunay sets $\mathcal{D}(T)$ are R-relatively dense. Let ν be the constant given by Lemma 3.13 for $\lambda = 3R$, so that if $v \in \mathcal{V}_{3R}$ and $T_1, T_2 \in \Xi_v$ agree up to radius ν^{-1}, then $h_{T_i}(v)$ does not depend on i.

Let $\delta > 0$, which is assumed to be small with respect to the radius of the tiles, and smaller than ν. Assume also that $\delta^{-1} > 3R$. By uniform continuity of h^{-1}[2], there is $\varepsilon > 0$ such that

$$\forall T_1, T_2 \in \Xi,\ D(T_1, T_2) \geq (\delta^{-1} + 3R)^{-1} \Rightarrow D(h(T_1), h(T_2)) \geq \varepsilon.$$

Let now $T_1, T_2 \in \Xi$, such that

$$D(T_1, T_2) \geq (\delta^{-1} + r)^{-1},$$

which means that these two tilings fail to agree on a ball of radius $r + \delta^{-1}$. Then, there exists $x \in \mathbb{R}^d$ with $\|x\| \leq r$, such that $T_1 - x$ and $T_2 - x$ fail to agree on a ball of radius δ^{-1}. The natural follow-up is to claim that $h(T_1) - h_{T_1}(x)$ and $h(T_2) - h_{T_2}(x)$ disagree within radius ε^{-1}, and since $\|h_{T_i}(x)\| \leq M \|x\|$, then $h(T_1)$ and $h(T_2)$ disagree within radius $Mr + \varepsilon^{-1}$. This line of reasoning fails because there is no guarantee that $h_{T_1}(x) = h_{T_2}(x)$. Some precautions have to be taken.

Run the following algorithm, which is illustrated by Figure 6:

1. Initialize $k = 0$, and start with $x_k = x_0 = 0$;

2. If $T_1 - x_k$ and $T_2 - x_k$ do not agree on $B(0, 3R + \delta^{-1})$, stop the algorithm. If they do agree (which means that $\|x - x_k\|$ must be at least $3R$), continue to step (3).

3. Using Lemma 3.11, there exists $x_{k+1} \in \mathcal{D}(T_i)$ such that $\|x_k - x_{k+1}\| \leq 3R$, and $\|x - x_{k+1}\| \leq \|x - x_k\| - R$. Since $T_1 - x_k$ and $T_2 - x_k$ agree at least up to radius $3R$, the choice of x_k can be made the same for T_1 and T_2.

---

[2] In its contrapositive form.

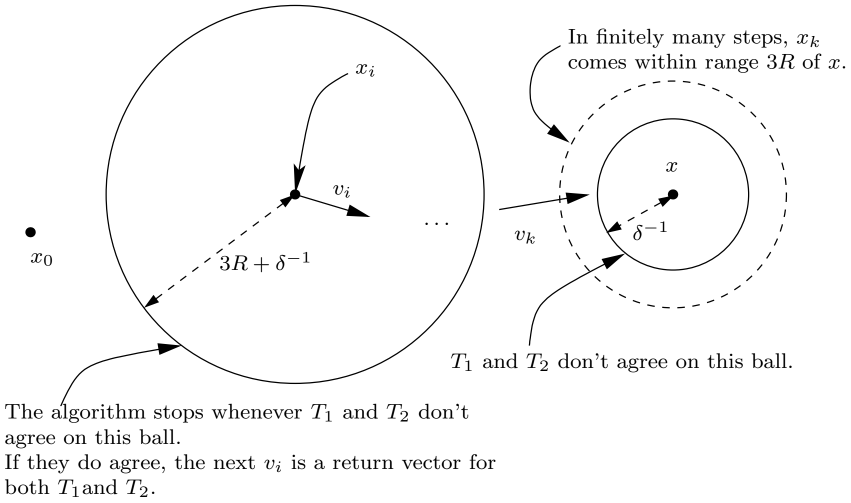
In finitely many steps, x_k comes within range $3R$ of x.
x_i
x
v_i
$\dots$
v_k
δ^{-1}
x_0
$3R+\delta^{-1}$
T_1 and T_2 don't agree on this ball.
The algorithm stops whenever T_1 and T_2 don't agree on this ball.
If they do agree, the next v_i is a return vector for both T_1 and T_2.

Figure 6: An illustration of the algorithm.

4. Return to step (2).

This algorithm terminates. Indeed, the distance between x_k and x is reduced at each step by at least R. After finitely many steps, x_k and x are within distance less than $3R$ of each other. When this happens, the stopping condition (2) applies, and the algorithm stops. The result of the algorithm is a finite sequence of vectors $(x_i)_{i=0,\dots,k}$ such that each $v_i := x_{i+1} - x_i$ is a return vector to Ξ in $\mathcal{V}_{3R}$. Furthermore, for all $i < k$, $T_1 - x_i$ and $T_2 - x_i$ agree up to radius at least δ^{-1}, which is greater than ν^{-1}. Therefore, by Lemma 3.13,

$$h_{T_1 - x_i}(v_i) = h_{T_2 - x_i}(v_i),$$

and these are return vectors to Ξ', which we note $(w_i)_{i=0,\dots,k-1}$.

Because of the stopping condition of the algorithm above,

$$D(T_1 - x_k, T_2 - x_k) \geq (3R + \delta^{-1})^{-1}, \text{ and therefore } D(h(T_1 - x_k), h(T_2 - x_k)) \geq \varepsilon.$$

Now, just remark that

$$\begin{aligned} h(T_i - x_k) &= h(T_i - v_0 - v_1 - \ldots - v_{k-1}) \\ &= h(T_i) - w_0 - \ldots - w_{k-1}. \end{aligned}$$

In particular, there is a *same* y such that $h(T_i - x) = h(T_i) - y$, $i = 1, 2$.

It was proved that $h(T_1) - y$ and $h(T_2) - y$ disagree on the ball of radius ε^{-1}. Additionally,

$$\begin{aligned} \|y\| &\leq M \|x_k\| \\ &\leq M(\|x\| + 3R), \end{aligned}$$

where the first inequality if obtained by Lemma 3.9, and the second by the fact that $\|x - x_k\| \leq 3R$.

To sum up, it was proved that whenever two tilings $T_1, T_2 \in \Xi$ disagree within radius $\delta^{-1} + r$, then their images have to disagree within radius at most $Mr + C$ for a constant $C = \varepsilon^{-1} + 3MR$. Therefore,

$$p_\Xi(r + \delta^{-1}) \leq p_{\Xi'}(Mr + C).$$

Up to changing M and restricting r to large values, one gets $p_\Xi(r) \leq p_{\Xi'}(Mr)$.

The other inequality is obtained by reversing the roles of h and h^{-1}. $\square$

The following result is an immediate consequence of this theorem.

Corollary 3.15. Let Ω be a (repetitive, aperiodic, FLC) tiling space, with polynomial complexity function, which means that there exists an α such that for some Ξ (and hence for all Ξ), there are two constants C_1 and C_2 such that for all r big enough,

$$C_1 r^\alpha \leq p(r) \leq C_2 r^\alpha.$$

Then the exponent α is a homeomorphism invariant of the space: any tiling space which is homeomorphic to Ω has polynomial complexity function with the same exponent α.

The main tool for showing that two homeomorphic tiling spaces have equivalent complexity functions is the study of how an orbit is sent to an orbit. A big part of the information carried by a homeomorphism lies in the ways orbits are deformed. How this relates to “tiling deformations” is addressed in the next section. Before moving to this topic, it is worth pointing out that the techniques which were applied to the complexity could also be applied to other quantities. Namely, it is reasonable to conjecture that the repetitivity function of a tiling[3] is preserved by homeomorphism (up to equivalence). An application of this conjecture would be that any tiling space which is homeomorphic to a linearly repetitive tiling space is itself linearly repetitive.

4 Interpretation in terms of deformations

There have been essentially two approaches for understanding tiling deformations. In the first one—due to Clark and Sadun [6]—a deformation is defined as a map from the set of edges of a polytopal tiling to $\mathbb{R}^d$. The image of an edge is a (a priori) different edge; fitting edges together, it defines a deformation of the tiling. This map needs to be *pattern-equivariant* in the sense that the image of an edge of T should be determined by the local pattern around it.

The other point of view uses Kellendonk's pattern-equivariant differential forms [9, 10]. A function defined on $\mathbb{R}^d$ is T-equivariant for some fixed tiling T if its value at x only depends on the local pattern of T around x. A deformation of a point pattern $\mathcal{D}(T)$ is then the image $f(\mathcal{D}(T))$ under a differentiable function $f : \mathbb{R}^d \to \mathbb{R}^d$, with T-equivariant differential.

---

[3] It is a function $r \mapsto M(r)$ which counts the size within which any patch of size r will repeat in a tiling.

In both cases, a deformation is associated with an element in the cohomology of the tiling space: the Čech cohomology on the one hand (interpreted in [6] by using the pattern-equivariant formalism), and the pattern-equivariant De Rham cohomology on the other hand. These cohomology groups turn out to agree [11, 6, 5].

The goal of this section is to explain how, starting with a homeomorphism $h : \Omega \to \Omega'$, one can associate a cohomology class in $\mathrm{H}^1(\Omega, \mathbb{R}^d)$. This class turns out to be unique modulo infinitesimals (or equivalently, it is unique in Kellendonk's mixed cohomology group). The formalism in this part is hybrid between Kellendonk's and Clark–Sadun's, and also uses some elements of the work of Savinien–Bellissard and Boulmezaoud–Kellendonk.

4.1 A few definitions

This section provides a brief overview of how deformations of tilings were defined, and how they are associated with cohomology. The Clark–Sadun point of view is used here.

Let T be a polygonal tiling, which is fixed from this point on. It gives $\mathbb{R}^d$ a CW-complex structure. Let $\mathcal{D}(T)$ be the set of vertices of T (the 0-skeleton). Assume $0 \in \mathcal{D}(T)$. Let $\mathcal{E}$ be its set of edges (it is formally a subset of the tangent space of the affine space $\mathbb{R}^d$, that is a subset of $\mathbb{R}^d \times \mathbb{R}^d$, where the first summand records the position of the starting point of the edge, and the second summand records the vector itself). Each edge is recorded twice with both orientations: if $(x, v) \in \mathcal{E}$, then $(x + v, -v) \in \mathcal{E}$. The notation is $-(x, v) = (x + v, -v)$.

The T-equivariant cellular cohomology is defined using the chain complexes of $\mathbb{R}^d$ viewed as a CW-complex, and the usual boundary operator. Strongly pattern-equivariant cochains valued in E are cochains which satisfy the following condition:

$$(T - x) \cap B(0, R) = (T - y) \cap B(0, R) \Longrightarrow \varphi(x, v) = \varphi(y, v),$$

for a given R, and for all (x, v) and all $(y, v) \in \mathcal{E}$. Weakly pattern-equivariant cochains f are such that for all $\varepsilon > 0$, there is a strongly equivariant cochain g_ε such that for all $e \in \mathcal{E}$,

$$\|f(e) - g_\varepsilon(e)\| < \varepsilon.$$

This cochain complex comes with the usual coboundary operator, and the coboundary of a strongly (resp. weakly) equivariant cochain is also strongly (resp. weakly) equivariant. The cohomology groups are denoted:

$$\mathrm{H}^k_{\mathrm{pe,s}}(\Omega_T; E) \quad \text{and} \quad \mathrm{H}^k_{\mathrm{pe,w}}(\Omega_T; E)$$

respectively.

As pointed out in [6], a strongly pattern-equivariant cochain is actually a cochain defined on the module of 1-chains of an appropriate finite CW-approximant of the space (a.k.a. Anderson–Putnam–Gähler complexes [1]). So strongly T-equivariant cellular cohomology is is isomorphic to Čech cohomology of Ω_T.

Let f and g be 1-cocycles, such that $f - g$ is the coboundary of a weakly pattern-equivariant 0-cochain. Then there is $s : \mathcal{D}(T) \to E$ such that for all edge $(x, v) \in \mathcal{E}$,

$$(f - g)(x, v) = s(x + v) - s(x).$$

Note that by weak pattern equivariance, this function s defines a continuous function $\tilde{s}$ on $\Xi \to E$, by extending the following uniformly continuous and densely defined function:

$$\tilde{s}(T - x) = s(x).$$

Equivariant cocycles with respect to T can always be interpreted more "globally", as objects defined on appropriate transversals of the tiling space. Here, this standpoint is used only for 0-cocycles, as it is the simplest; the reader can refer to the definition of PV-cohomology [19], and its comparison with pattern-equivariant cohomology in [5].

The relation $f \sim g$ if $f - g$ is the boundary of a weakly pattern equivariant function is an equivalence relation on cocycles, and defines an equivalence relation on strong cohomology classes. The quotient group is called the mixed group:

$$\mathrm{H}_{\mathrm{pe,m}} = \{\text{strongly equivariant 1-cocycles}\} / \sim .$$

It is also said that $f - g$ is *asymptotically negligible* if its image in the mixed group is zero. This is not the original definition of asymptotically negligible as it can be found in [6], but it is equivalent (see [10]).

Remark 4.1. Kellendonk's pattern equivariant De Rham cohomology allows to define analogous groups $\mathrm{H}^k_{\mathrm{DR}}$ using differential forms (this is how the weak and mixed group were first defined). Weak, strong, or mixed cohomology groups do not depend on the way they are defined, see [5].

$$\mathrm{H}^k_{\mathrm{DR},\bullet}(\Omega_T; E) \simeq \mathrm{H}^k_{\mathrm{pe},\bullet}(\Omega_T; E) \quad \text{for } \bullet \in \{\mathrm{s}, \mathrm{w}, \mathrm{m}\}.$$

Deformations of tilings are closely related to the cohomological objects above. The *data encoding a deformation* is a map $f : \mathcal{E} \to \mathbb{R}^d$ satisfying the following conditions:

- Strong pattern equivariance: there is an R such that for all $e = (x, v) \in \mathcal{E}$, the value of $f(x, v)$ only depends on $(T - x) \cap B(0, R)$;
- $f(-e) = -f(e)$;
- The image of a tile is a tile.

The last condition is actually a cocycle condition: a 2-dimensional tile (or face of a tile) consists of a closed cycle of edges. The condition that it is a face implies that the images of these edges have to sum to zero.

Given such a deformation data f, the deformation map F is defined as follows: if

$$x = \sum_{i=1}^{n} v_i \text{ with } \forall i, \ \left(\sum_{j=1}^{i-1} v_j, v_i \right) \in \mathcal{E},$$

then

$$F(x) = \sum_{i=1}^{n} f \left(\sum_{j=1}^{i-1} v_j, v_i \right).$$

In other word, $F(x)$ is computed by choosing a path of edges to x, and concatenating the images by f of the edges. Two paths from 0 to x differ by a boundary, so choosing two different paths gives the same image. Note that the deformation is defined on the 0-skeleton of the tiling T. It is therefore a map from a Delaunay set to an other point set.

Remark 4.2. At this point, it should be assumed that in order to be called a deformation, F sends $\mathcal{D}(T)$ to a Delaunay set. This is the case when f is a small deformation data (that is $f(x, v)$ is close to v). More generally, it seems sufficient to assume that f deforms any tile in a tile with non-empty interior and either consistently preserves or consistently changes the orientation. It is unknown to the author whether there is a necessary and sufficient condition which can be stated on the deformation data f, and is equivalent to the fact that $F(\mathcal{D}(T))$ is Delaunay.

Assume f and g are two deformation data such that $f - g$ is a coboundary. There is then a T-equivariant function s defined on $\mathcal{D}(T)$ such that $f - g = \delta s$. A computation shows that $F - G = s - s(0)$. If one sees s as a function defined on Ξ, it is equivalent to $(F - G)(x) = s(T - x) - s(T)$.

4.2 Tiling space homeomorphisms as deformations

In light of the discussion above, it appears that tiling space homeomorphisms are related to deformations. Let $h : \Omega \to \Omega'$ be a homeomorphism, and h_ε, Ξ, Ξ' be the data provided by Lemma 3.8. Let $\mathcal{D}$ and $\mathcal{D}'$ be the local pointing rules associated with Ξ and Ξ', and let $T \in \Xi$.

Let $\mathcal{T}$ be the local derivation introduced in Lemma 3.9, such that for $T \in \Xi$, $\mathcal{T}(T)$ is a triangulation with vertices $\mathcal{D}(T)$. Let $\mathcal{T}'$ be the image triangulation rule, with vertices given by the pointing rule $\mathcal{D}'$, and such that two vertices are linked by an edge in $\mathcal{T}'$ if and only if their preimages by h_ε are linked by an edge in $\mathcal{T}$. The image of the triangulation does not need to be a triangulation itself (see Figure 5).

Then h_ε induces a deformation. Indeed, let v be an edge of the triangulation $\mathcal{T}$, and $\Xi_v := \Xi \cap \Xi - v$. Then the map $T \mapsto h_{\varepsilon,T}(v)$ is continuous, and its image is a return vector to Ξ'. By finite local complexity of $\mathcal{D}'(T')$ ($T' \in \Omega'$), the range of this map is finite, and so it is locally constant. Therefore, the is an $R_v > 0$ such that $h_{\varepsilon,T}(v)$ only depends on $T \cap B(0, R_v)$ (see also Lemma 3.13). Taking $R = \max\{R_v\}$, there is a strongly pattern-equivariant deformation data f with radius R. For all $T \in \Xi$, f induces a deformation $F_T : \mathcal{T}(T) \mapsto \mathcal{T}'(h_\varepsilon(T))$. This function F_T coincides with h_ε on Ξ, and is the function g defined in Lemma 3.9.

Because of this, h is associated with a cohomology class in $\mathrm{H}^1_{\mathrm{s}}(\Omega; \mathbb{R}^d)$. This class depends a priori of the choice of Ξ, Ξ' and h_ε. However, by Lemma 3.8, there is a function s such that $h_\varepsilon = h - s$ for a continuous s defined on Ω. In the "leaf-wise" picture, it amounts to:

$$h_{\varepsilon,T}(x) = h_T(x) - s(T - x).$$

In other words, h_ε and h differ by a weakly equivariant coboundary. Therefore, two different choices of Ξ, Ξ' and h_ε would lead to deformations which differ by a coboundary. Therefore, the cohomology class of h is defined unambiguously in the mixed group.

Proposition 4.3. *Let $h : \Omega \longrightarrow \Omega'$. Then, there is a well-defined cohomology class $[h]_m$ in the mixed cohomology group, such that*

$$[h]_m := [h_\varepsilon]_m$$

for any deformation h_ε given by Lemma 3.8. The image does not depend on the choice of h_ε.

In addition, whenever two homeomorphisms correspond to the same class (i.e. they differ by a coboundary), their images are topologically conjugate, as stated by the following proposition.

Proposition 4.4. *Let $h_1 : \Omega \longrightarrow \Omega'$ and $h_2 : \Omega \longrightarrow \Omega''$ be homeomorphisms, which differ by a weak co-boundary, that is there is a continuous function $s : \Omega \to \mathbb{R}^d$ such that for all $T \in \Omega$ and all $x \in \mathbb{R}^d$,*

$$h_{2,T}(x) = h_{1,T}(x) - (s(T) - s(T - x)). \tag{3}$$

Then the tiling spaces Ω' and Ω'' are topologically conjugate.

Proof. The proof is essentially an adaptation of the one in [6], with a slightly different formalism. A preliminary computation shows that for a homeomorphism h between tiling spaces, the following homeomorphisms induced on the leaves are equal: $(h^{-1})_{h(T)} = (h_T)^{-1}$ (this is left as an exercise).

For all $T' \in \Omega'$ and all $y \in \mathbb{R}^d$, let us compute $h_2 \circ h_1^{-1}(T' - y)$. Let $T = h_1^{-1}(T')$.

$$\begin{aligned}
h_2 \circ h_1^{-1}(T' - y) &= h_2\big(T - h_{1,T}^{-1}(y)\big) \\
&= h_2(T) - h_{2,T} \circ h_{1,T}^{-1}(y) \\
&= h_2(T) - h_{1,T} \circ h_{1,T}^{-1}(y) - s\big(T - h_{1,T}^{-1}(y)\big) + s(T) \\
&= h_2 \circ h_1^{-1}(T') - y - s\big(h_1^{-1}(T') - (h_1^{-1})_{T'}(y)\big) + s\big(h_1^{-1}(T')\big) \\
&= h_2 \circ h_1^{-1}(T') - y - s \circ h_1^{-1}(T' - y) + s \circ h_1^{-1}(T').
\end{aligned}$$

Now, given $T' \in \Omega'$, it is possible to define

$$\varphi(T') = h_2 \circ h_1^{-1}(T') - s \circ h_1^{-1}(T').$$

By the computation above, it commutes to the action, and is therefore a bijection leaf-wise. Since h_1 and h_2 are bijective, φ induces a bijection between the set of leaves $\Omega'/\mathbb{R}^d \to \Omega''/\mathbb{R}^d$. Putting things together, φ is bijective. It is obviously continuous. It is therefore a homeomorphism by compactness. In the end, it is a topological conjugacy. □

Conversely, from a recent work of Kellendonk and Sadun [12], a topological conjugacy is (up to mutual local derivations) essentially a weakly pattern-equivariant deformation. It seems to provide a converse of Proposition 4.4: whenever the following commutative triangle occurs,

$$\begin{tikzcd}
\Omega \arrow[r, "h_1"] \arrow[dr, "h_2"'] & \Omega' \arrow[d, "\varphi"] \\
 & \Omega''
\end{tikzcd}$$

with h_i homeomorphisms, and φ a topological conjugacy, then h_1 and h_2 differ by a weak coboundary, and therefore have the same cohomology class in the mixed group.

As a conclusion, it seems that the mixed cohomology group of a tiling space Ω can be used to describe homeomorphisms to (generally) other tiling spaces up to topological conjugacy. More precisely, the cohomology class describes how the homeomorphism h deforms the action of $\mathbb{R}^d$ on the leaves. It is interesting to point out that in the finite local complexity case, there is always a strongly pattern-equivariant deformation underlying a homeomorphism (up to a weak co-boundary): it is not the weak cohomology group which is used.

4.3 Discussion on a groupoid interpretation

As remarked in [5], pattern-equivariant cohomology theory can be approached using groupoid cohomology. It turns out that deformations can be described quite simply using the groupoid, and the link with cohomology is then quite direct.

Following [17], given a topological groupoid $\mathcal{G}$, a 1-cocycle with values in a group G (say $G = \mathbb{R}^k$ or $\mathbb{Z}^k$) is a groupoid morphism $\mathcal{G} \to G$ (where G is seen as a groupoid with a unique base point). A 1-coboundary is a morphism which is induced by a continuous function $b : \mathcal{G}^{(0)} \to G$ as follows:

$$\varphi_b(x) = b(r(x)) - b(s(x)),$$

where r and s are the range and source maps, respectively.

Clark–Sadun's approach in the groupoid setting Let $\Xi \subset \Omega$ be a transversal and $\mathcal{V}$ be the set of return vectors to this transversal (Definition 3.12). Let

$$\mathcal{G} = \{(T, v) \in \Xi \times \mathcal{V}\,;\, T - v \in \Xi\}$$

be the reduction of the tiling groupoid to a transversal. By a slight abuse of terminology, this groupoid is noted $\Xi \rtimes \mathcal{V}$. Then a morphism $\mathcal{G} \to \mathbb{R}^d$ is a map which satisfies:

$$f((T - x, y)(T, x)) = f(T - x, y) + f(T, x),$$

or equivalently

$$f(T, x + y) - f(T - x, y) - f(T, x) = 0.$$

It can be restated as follows: whenever there exists a tiling T in which there is a cycle of return vectors x, y, $-(x + y)$ which sum to zero, their image by f is zero. It is the cocycle condition which is also required to define deformation data. In this formalism, given a triangulation with edge set $\mathcal{E}$, a deformation data is a map

$$f : \{(T, v) \in \mathcal{G}\,;\, v \in \mathcal{E}\} \longrightarrow \mathbb{R}^d$$

which is locally constant, and can be extended to a groupoid morphism $\tilde{f} : \mathcal{G} \to \mathbb{R}^d$. Note that the subset on which f is defined is a generating set for $\mathcal{G}$. Then for a give $T \in \Xi$, the deformation induced by f is then given by this map defined on $\mathcal{D}(T)$:

$$F_T(x) = f(T, x) \quad \text{for } x \in \mathcal{D}(T).$$

In the end, a deformation data gives rise to a groupoid morphism, which is itself a cocycle.

Kellendonk's approach in the groupoid setting In the Kellendonk framework the appropriate groupoid is $\Omega \rtimes \mathbb{R}^d$. A deformation F is a map $\mathbb{R}^d \to \mathbb{R}^d$, such that the differential df is T-equivariant. Using the equation

$$F(x + h) - F(x) = \int_0^1 dF(x + th) \cdot h \, dt,$$

one proves that for all fixed h, the function $x \mapsto F(x + h) - F(x)$ is pattern-equivariant. Therefore, it defines a groupoid morphism $\Omega \rtimes \mathbb{R}^d \to \mathbb{R}^d$, which is the extension by continuity of the following map:

$$(T - x, h) \longmapsto F(x + h) - F(x).$$

Therefore, it defines a $\mathbb{R}^d$-valued cohomology class of the groupoid $\Omega \rtimes \mathbb{R}^d$.

Tiling space homeomorphisms in the groupoid setting Last, consider a homeomorphism $h : \Omega \longrightarrow \Omega'$. Then, by Theorem 3.4, the map

$$(T, x) \longmapsto h_T(x)$$

is a continuous map, and a direct check proves that it is a groupoid morphism. Therefore, it is a cocycle, and represents a cohomology class.

It appears that groupoids provide a convenient common framework for addressing both deformations and cohomology of tiling spaces. As it appears, homeomorphisms between tiling spaces also fall in this framework.

However, attention should be paid to several points: the groupoid cohomology could be represented by continuous, or locally constant cocycles in the case of the cohomology of the groupoid $\Xi \rtimes \mathcal{V}$. It would lead to different cohomology groups in general (when the ring of coefficients is $\mathbb{R}$, for example). It is possible to define an analogue of the strongly equivariant and weakly equivariant cohomology group, as well as a mixed group.

Similarly when the groupoid is $\Omega \rtimes \mathbb{R}^d$, it is possible to use continuous cocycles, transversally locally constant cocycles, differentiable cocycles. It would lead a priori to different cohomology theories. Some of these groupoid cohomology theories were compared in [5]; one can also refer to work by Kellendonk and Putnam [11], where they study the cohomology of the group $\mathbb{R}^d$ with values in the module $C^\infty(\Omega; \mathbb{R})$ of functions which are differentiable along the leaves and possibly transversally locally constant. This "dynamic cohomology" is related with the cohomology of the groupoid.

Even if some of the formal work of verification has yet to be done, it is reasonable to conjecture that, regardless of the groupoid which is being considered, one can form weakly pattern-equivariant cohomology groups, strongly pattern-equivariant groups and mixed groups, and that these groups are isomorphic to the eponymous previously defined groups.

4.4 Concluding remarks

It was proved that whenever two aperiodic repetitive tiling spaces with finite local complexity are homeomorphic, the homeomorphism is not far from being Lipschitz along the leaves (Lemma 3.9), and therefore has to preserve the complexity function, to some extent. Such a homeomorphism gives rise to a cohomology class in a mixed group, which parametrizes deformations up to topological conjugacy. Some questions still remain.

- It would be tempting to state that a tiling space homeomorphism "is" a deformation, up to topological conjugacy. However, this might be a bit fast. Given $T \in \Omega$, it is possible to derive $\mathcal{D}(T)$ and to deform into $\mathcal{D}'(h(T))$, but this second point set does not need to be MLD to T', it is just locally derived (see for example Figure 1). It could be possible to fix this problem by finding a sub-transversal $\Xi'' \subset \Xi'$ such that the new (sparser) Delaunay set $\mathcal{D}''(h(T))$ is MLD to $h(T)$. While this might be possible for a large family of tilings (including those which are uniquely ergodic and have zero entropy), it is unknown to the author whether it can always be done.

- In relation with the previous point: since the point pattern $\mathcal{D}'(h(T))$ does not need to be MLD to $h(T)$, its tiling space could be a strict factor of Ω'. In other words, the deformation underlying the homeomorphism $h : \Omega \to \Omega'$ could also be a deformation underlying a factor map. Therefore, the mixed group may represent deformations up to topological conjugacy, but probably not homeomorphisms up to topological conjugacy. It may only be true in frameworks where the maps are *assumed* to be homeomorphisms.

- Given a homeomorphism $\Omega \to \Omega'$ which sends a transversal Ξ to a transversal Ξ', it is possible to find a strongly pattern-equivariant deformation data which induces a deformation of the pointing rule $\mathcal{D}(T)$ into $\mathcal{D}'(T)$ for $T \in \Xi$, as explained in this section (it is also the map g of Lemma 3.9). It was also proved by Kellendonk in [10, Theorem 4.10]. However, is it possible to ensure that the deformation maps $\mathbb{R}^d \to \mathbb{R}^d$ are homeomorphisms? In other words, given a homeomorphism $h : \Omega \to \Omega'$, is it possible to find another homeomorphism $h' : \Omega \to \Omega'$, which differs from h by a weak cocycle (in the sense of Proposition 4.4), and which is induced by a strongly pattern equivariant deformation data? Equivalently, such that the induced groupoid morphism $(T, x) \mapsto h'_T(x)$ is locally constant in the first variable. I would conjecture that the answer is yes, but the topological technicalities might be quite important.

References

[1] J. E. Anderson and I. F. Putnam. Topological invariants for substitution tilings and their associated C^*-algebras. *Ergodic Theory Dynam. Systems*, 18(3):509–537, 1998.

[2] M. Baake, M. Schlottmann, and P. D. Jarvis. Quasiperiodic tilings with tenfold symmetry and equivalence with respect to local derivability. *J. Phys. A*, 24(19):4637–4654, 1991.

[3] J. Bellissard, R. Benedetti, and J.-M. Gambaudo. Spaces of tilings, finite telescopic approximations and gap-labeling. *Comm. Math. Phys.*, 261(1):1–41, 2006.

[4] R. Benedetti and J.-M. Gambaudo. On the dynamics of $\mathbb{G}$-solenoids. Applications to Delone sets. *Ergodic Theory Dynam. Systems*, 23(3):673–691, 2003.

[5] H. Boulmezaoud and J. Kellendonk. Comparing different versions of tiling cohomology. *Topology Appl.*, 157(14):2225–2239, 2010.

[6] A. Clark and L. Sadun. When shape matters: deformations of tiling spaces. *Ergodic Theory Dynam. Systems*, 26(1):69–86, 2006.

[7] A. Julien. *Complexité des pavages apériodiques : calculs et interprétations.* PhD thesis, Université Lyon 1, Dec. 2009.

[8] A. Julien. Complexity and cohomology for cut-and-projection tilings. *Ergodic Theory Dynam. Systems*, 30(2):489–523, 2010.

[9] J. Kellendonk. Pattern-equivariant functions and cohomology. *J. Phys. A*, 36(21):5765–5772, 2003.

[10] J. Kellendonk. Pattern equivariant functions, deformations and equivalence of tiling spaces. *Ergodic Theory Dynam. Systems*, 28(4):1153–1176, 2008.

[11] J. Kellendonk and I. F. Putnam. The Ruelle-Sullivan map for actions of $\mathbb{R}^n$. *Math. Ann.*, 334(3):693–711, 2006.

[12] J. Kellendonk and L. Sadun. Meyer sets, topological eigenvalues, and cantor fiber bundles. *Preprint*, 2012.

[13] J. C. Lagarias and P. A. B. Pleasants. Repetitive Delone sets and quasicrystals. *Ergodic Theory Dynam. Systems*, 23(3):831–867, 2003.

[14] M. Morse and G. A. Hedlund. Symbolic Dynamics. *Amer. J. Math.*, 60(4):815–866, 1938.

[15] P. S. Muhly, J. N. Renault, and D. P. Williams. Equivalence and isomorphism for groupoid C^*-algebras. *J. Operator Theory*, 17(1):3–22, 1987.

[16] N. Priebe Frank and L. Sadun. Topology of some tiling spaces without finite local complexity. *Discrete Contin. Dyn. Syst.*, 23(3):847–865, 2009.

[17] J. Renault. *A groupoid approach to C^*-algebras*, volume 793 of *Lecture Notes in Mathematics.* Springer, Berlin, 1980.

[18] L. Sadun and R. F. Williams. Tiling spaces are Cantor set fiber bundles. *Ergodic Theory Dynam. Systems*, 23(1):307–316, 2003.

[19] J. Savinien and J. Bellissard. A spectral sequence for the K-theory of tiling spaces. *Ergodic Theory Dynam. Systems*, 29(3):997–1031, 2009.

[20] M. Schlottmann. Periodic and quasi-periodic Laguerre tilings. *Internat. J. Modern Phys. B*, 7(6-7):1351–1363, 1993.